Relevance definition rubric

A practical rubric to define relevance, align teams, and prevent “AI magic” ambiguity.

PURPOSE

Define “relevance” in a way that product, design, and data science can all execute against—without ambiguity.

Step 1: Define the job-to-be-done

- ■ Primary user job is stated in plain language (not a feature).
- ■ We know the moment: what triggers intent now?
- ■ We identified the top 3 user questions in this moment.
- ■ We know what “good outcome” looks like for the user.

Step 2: Specify signals and constraints

Signal type	Examples	Constraints / notes
Behavioral	Clicks, searches, tool usage, dwell time, repeats	Prefer explicit intent; avoid over-fitting to noise.
Account / context	Lifecycle stage, balances, plan type, tenure	Document sensitive fields; ensure policy/consent.
Content metadata	Topic, complexity, format, freshness, suitability	Keep taxonomy consistent; avoid “misc” tags.
Negative signals	Dismissals, skips, low dwell, exits	Use to reduce repetition and annoyance.

Relevance rubric (scoring)

Score each candidate item 1-5 in each dimension. Ship only if it clears your minimum bar (example: >= 16/25) and passes guardrails.

Dimension	1 (poor)	3 (ok)	5 (excellent)
Intent match	Unrelated to current goal	Somewhat related	Directly supports next best step
User value	Nice-to-have / vague	Helpful for some	Meaningfully reduces effort or improves decision
Timing	Wrong moment	Acceptable timing	Perfect moment based on signals
Trust / explainability	Hard to justify	Explainable with effort	Easy to explain in one sentence
Safety / appropriateness	Risky / sensitive	Low risk with guardrails	Clearly safe and appropriate

Definition of done

- ■ We documented the minimum bar to show an item (score threshold).
- ■ We documented exclusion rules (what must never be recommended).
- ■ We documented what the “Why this?” explanation will say.
- ■ We documented how we’ll evaluate relevance post-launch (metrics + feedback).